

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

PETITION FEE Under 37 CFR 1.17(f), (g) & (h) TRANSMITTAL (Fees are subject to annual revision)		Application Number	10/823,619
Send completed form to: Commissioner for Patents P.O. Box 1450, Alexandria, VA 22313-1450		Filing Date	April 14, 2004
		First Named Inventor	T. AMANO, et al
		Art Unit	
		Examiner Name	
		Attorney Docket Number	274.43199X00

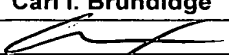
Enclosed is a petition filed under 37 CFR §1.102(d) that requires a processing fee (37 CFR 1.17(f), (g), or (h)). Payment of \$ 130.00 is enclosed.

This form should be included with the above-mentioned petition and faxed or mailed to the Office using the appropriate Mail Stop (e.g., Mail Stop Petition), if applicable. For transmittal of processing fees under 37 CFR 1.17(i), see form PTO/SB/17i.

Payment of Fees (small entity amounts are NOT available for the petition (fees))

- ☒ The Commissioner is hereby authorized to charge the following fees to Deposit Account No. 50-1417:
- ☐ petition fee under 37 CFR 1.17(f), (g) or (h) ☒ any deficiency of fees and credit of any overpayments
- Enclose a duplicative copy of this form for fee processing.
- ☐ Check in the amount of \$ _____ is enclosed.
- ☒ Payment by credit card (From PTO-2038 or equivalent enclosed). Do not provide credit card information on this form.

Petition Fees under 37 CFR 1.17(f): For petitions filed under: § 1.53(e) - to accord a filing date. § 1.57(a) - to according a filing date. § 1.182 - for decision on a question not specifically provided for. § 1.183 - to suspend the rules. § 1.378(e) for reconsideration of decision on petition refusing to accept delayed payment of maintenance fee in an expired patent. § 1.741(b) - to accord a filing date to an application under §1.740 for extension of a patent term.	Fee \$400	Fee Code 1462
Petition Fees under 37 CFR 1.17(g): For petitions filed under: §1.12 - for access to an assignment record. §1.14 - for access to an application. §1.47 - for filing by other than all the inventors or a person not the inventor. §1.59 - for expungement of information. §1.103(a) - to suspend action in an application. §1.136(b) - for review of a request for extension of time when the provisions of section 1.136(a) are not available. §1.295 - for review of refusal to publish a statutory invention registration. §1.296 - to withdraw a request for publication of a statutory invention registration filed on or after the date the notice of intent to publish issued. §1.377 - for review of decision refusing to accept and record payment of a maintenance fee filed prior to expiration of a patent. §1.550(c) - for patent owner requests for extension of time in <u>ex parte</u> reexamination proceedings. §1.956 - for patent owner requests for extension of time in <u>inter partes</u> reexamination proceedings. § 5.12 - for expedited handling of a foreign filing license. § 5.15 - for changing the scope of a license. § 5.25 - for retroactive license.	Fee \$200	Fee code 1463
Petition Fees under 37 CFR 1.17(h): For petitions filed under: §1.19(g) - to request documents in a form other than that provided in this part. §1.84 - for accepting color drawings or photographs. §1.91 - for entry of a model or exhibit. §1.102(d) - to make an application special. §1.138(c) - to expressly abandon an application to avoid publication. §1.313 - to withdraw an application from issue. §1.314 - to defer issuance of a patent.	Fee \$130	Fee Code 1464

Name (Print/Type)	Carl I. Brundidge	Registration No. (Attorney/Agent)	29,621
Signature		Date	August 4, 2005

This collection of information is required by 37 CFR 1.114. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.



274.43199X00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: T. AMANO, et al

Serial No.: 10/823,619

Filed: April 14, 2004

For: METHOD AND APPARATUS FOR AVOIDING JOURNAL
OVERFLOW ON BACKUP AND RECOVERY SYSTEM USING
STORAGE BASED JOURNALING

PETITION TO MAKE SPECIAL
UNDER 37 CFR §1.102(MPEP §708.02)

MS Petition

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

August 4, 2005

Sir:

Applicants hereby petition the Commissioner to make the above-identified application special in accordance with 37 CFR §1.102(d). Pursuant to MPEP §708.02(VIII), Applicants state the following.

(A) This Petition is accompanied by the fee set forth in 37 CFR §1.17(h).

The Commissioner is hereby authorized to charge any additional payment due, or to credit any overpayment, to Deposit Account No. 50-1417.

(B) All claims are directed to a single invention.

If the Office determines that all claims are not directed to a single invention, Applicant will make an election without traverse as a prerequisite to the grant of special status in conformity with established telephone restriction practice.

08/05/2005 JBALINAN 00000043 10823619

01 FC:1464

130.00 OP

(C) A pre-examination search has been conducted.

The search was directed towards a backup and recovery method as set forth in claims 1-23, particularly claims 1, 6, 12 and 19. According to the present invention, the backup and recovery method for a storage system that avoids journal overflow and includes producing at least a first snapshot of a data volume which is configured to receive data by way of write operations issued from a host device; producing a journal entry for each write operation issued from the host device; storing each journal entry in a journal volume, thereby accumulating a list of journal entries; monitoring an amount of free space on the journal volume; and when the free space falls below a threshold value, taking a new snapshot of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow. Under additional aspects, the present invention provides that the threshold value is set in a management table that includes a plurality of entries each containing information of respective journal volumes and that the threshold value is an indication of the lowest amount of free capacity of the journal pool the storage system is allowed to reach.

The search of the above features was conducted in the following areas:

<u>Class</u>	<u>Subclasses</u>
707	200, 202, 203, 204
711	114, 147, 152, 159, 161, 162
714	2, 13

Additionally, a computer database search was conducted on the USPTO systems EAST and WEST.

(D) The following is a list of the references deemed most closely related to the subject matter encompassed by the claims:

<u>U.S. Patent Number</u>	<u>Inventors</u>
4,077,059	Cordi et al
4,823,261	Bank et al
6,128,630	Shackelford
6,189,016	Cabrera et al
6,324,654	Wahl et al
6,442,706	Wahl et al
<u>Foreign Documents</u>	<u>Inventor</u>
JP 3103941	Kariya et al
JP 5002517	Abe
JP 11353215	Hata
JP 2000155708	Kondo

A copy of each of these references (as well as other references uncovered during the search) is enclosed in an accompanying IDS.

(E) It is submitted that the present invention is patentable over the references for the following reasons.

It is submitted that the cited references, whether taken individually or in combination with each other, fail to teach or suggest the invention as claimed. In particular, the cited references, at a minimum, fail to teach or suggest as recited in the claims:

a first feature of the present invention as recited in independent claim 1 wherein when the free space falls below a threshold value, taking a new snapshot of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow,

a second feature of the present invention as in independent claim 6, wherein when the free space falls below a threshold value, taking a logical snapshot of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow,

a third feature of the present invention as recited in independent claim 12, wherein when the free space falls below a threshold value, stopping the storing of journal entries, switching to bit map management and taking a logical snapshot of the data volume, thereby avoiding journal overflow, and

a fourth feature of the present invention as recited in independent claim 19, wherein when the free space falls below a threshold value, overwriting the oldest journal entry stored on the journal volume, thereby avoiding journal overflow.

Further, the cited references fail to teach or suggest the above noted features of the present invention when taken in combination with the other limitations recited in the claims.

The references considered most closely related to the claimed invention are briefly discussed below:

Cordi (U.S. Patent No. 4,077,059) discloses a hierarchical memory system for a multi-processing system which has two or more processing units accessing the memory system. The memory system has two different types of memory units on each level. One of the types of units is called the data store (DS) and contains all the data at that level of the memory. The other type of unit is called the copy back data store (CBDS) and contains all the changes that have been

made which changes data either by addition or modification and that are to be copied back to the next lower level of the memory hierarchy. The data store and the copy back data store in each level are on two different power systems and transfers of the changes to the next lower level are done in the order in which the change entered in the copy back store with the oldest entry being the first to be copied back. Each copy back data store has a capacity which is only a small portion of the data capable of being stored in the corresponding data store. To prevent overruns from occurring when the number of data units to be copied back exceeds the capacity of the copyback store, an up/down counter keeps track of the difference between the number of items to be copied back and those already copied back and notifies the system when the copy back data store is about to overflow. A journal full counter is used to ensure that the next free space in counter does not overrun the copy back counter and start overlying entries in the journal on CBDS that have not as of yet been copied back. The journal full register is an up/down counter that keeps count of the number of the m slots in the journal and the CBDS that are used and not copied back. When a certain percentage of these slots fall into this category the journal full counter issues a regenerate copy back command to the circuitry of FIG. 6 to initiate a copy back operation. For instance, it may be advantageous to issue copy back requests when the journal reaches a threshold of 90% capacity and continue copy back cycles until 60% capacity is reached. (See, e.g., Abstract and column 8, lines 26-39).

However, unlike the present invention, Cordi does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Cordi at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Bank (U.S. Patent No. 4,823,261) discloses an apparatus and method that employs dual checkpoint data sets for communicating system status. A journal of changed data is implemented to reduce I/O to a subsystem's shared data area on a non-volatile shared storage device. The journal provides for an increase in the amount of time that a processor may have access to the shared data area. Also, two versions of the data area are implemented in order to insure the integrity of the continuously updated data area. The two versions flip-flop depending upon which one has the most recent updates. That is, the version

that has the most recent updates becomes the to-be-read-from data area and is read by the processor that currently has access to the shared data area during this series of I/O operations. The other version becomes the to-be-written-to data area and is written to by the processor that currently has access to the data area in order to update the to-be-written-to version to the current level. The to-be-written-to version then becomes the to-be-read-from version during the next series of I/O operations. If the journal space is exhausted, (called a journal overflow), then the record is updated directly and all eighth bits are turned on. Control blocks can span multiple journal records. If the updated control blocks cannot be contained within the journal records, this is known as a journal overflow. There is code to handle this condition, should it occur. Essentially, updated control blocks that cannot be moved to the journal because it is full, will be written in their 4K pages (as in the prior art). (See, e.g., Abstract, column 17, lines 64-68, and column 20, lines 37-44).

However, unlike the present invention, Bank does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Bank at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1,

the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Kariya (Japanese Patent No. 3103941) discloses a method to simplify a program by providing a before journal monitor means, a lock list monitor means, and a roll back monitor means to eliminate a need for a user to preliminarily forecast the volume of a before journal and a lock list. A before journal monitor means which performs the commitment processing to abandon the before journal and the lock list of a process at the time of forecasting overflow of the before journal and a lock list monitor means which performs the commitment processing at the time of forecasting overflow of the lock list are provided. A roll back monitor means is provided which performs the commitment processing after execution of the roll back processing caused by dead lock after the last commitment processing of the process. Thus, it is unnecessary for a user to pay his attention to the timing of the commitment processing, and the program is simplified, and the before journal and the lock list are effectively used at its maximum to improve the processing performance. (See, e.g., Abstract).

However, unlike the present invention, Kariya does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free

space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Kariya at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Abe (Japanese Patent No. 5002517) discloses a method for safely using a data base system over a long period by surely executing the switching of journal files for writing data base journal records. At the time of detecting overflow when a journal record output means is writing records in a previously set journal file #1, a journal file automatic swapping means switches the working journal file #1 to another journal file #2. When a forced switching instruction is applied based upon a command outputted from a console, the working file #1 is forcedly switched by a journal file forced swapping means. Immediately before the switching, the writing file #1 is saved to a magnetic file by a journal file saving means and the file #1 is released. (See, e.g., Abstract).

However, unlike the present invention, Abe does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Abe at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Hata (Japanese Patent No. 11-353215) discloses a method for decreasing a journal-after-update quantity more by judging how update areas overlap and updating update information according to the judgment result. An update area overlap judging means judges whether the update area indicated by new update information overlaps with update areas indicated by respective pieces of update information in an update information storage area, an update information storage means adds the new update information to the update information storage area, and a start position change means overwrites the start position of the new update

information to the start position of the update information in an update information storage area where the new update area and an area overlap with each other; and an area length varying means overwrites new area length to the area length of the update information in the update information storage area showing the update area overlapping with the new update area and an update information integrating means judges whether or not the changed update information overlaps with the existent update information, and then change the start position or vary the area length of the existing update information and initializes one piece of overlapping update information when the changed update information overlaps. (See, e.g., Abstract).

However, unlike the present invention, Hata does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Hata at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach

or suggest these features in combination with the other limitations recited in each of the independent claims.

Kondo (Japanese Patent No. 2000-155708) discloses a method for preventing an overflow of a journal file from causing troubles to the whole system. A monitor means always monitor the capacities of a rollback journal file and a history journal file. When the use rate of each file exceeds a certain value, a console alarms. Then a program exceeding the use quantity of the rollback journal file is stopped. The saving is performed for the history journal file. (See, e.g., Abstract).

However, unlike the present invention, Kondo does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Kondo at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach

or suggest these features in combination with the other limitations recited in each of the independent claims.

Shackelford (U.S. Patent No. 6,128,630) discloses a log-structured storage for use with journal data sets and a two-phase commit process. The log-structured storage provides a space release function for the journal data sets, wherein one or more portions of the journal data sets that are no longer required are marked as free space. The log-structured storage of the present invention provides a space release function for journal data sets, which the application uses to mark the outdated data in the journal data set as being free space. When a cleaning process is executed for the log-structured storage, it does not have to move data in the journal data set unnecessarily, thereby resulting in a significant improvement in performance. Moreover, the application controls the invocation of the space release function, which provides a significant performance benefit over existing functions. See, e.g., Abstract and column 2, lines 38-47).

However, unlike the present invention, Shackelford does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Shackelford at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Cabrera (U.S. Patent No. 6,189,016) discloses a change journal for recording changes to files in a storage volume of a computer system keeps a record for each notable change to a file. Each record is given a unique update sequence number, which is a serial number of ever-increasing value assigned to each record. Each record includes one or more change reasons specifying what type of action occurred with respect to the associated file. The presence of a close file change reason in a record connotes that the record includes all the notable changes made to the file in a preceding file session, which is defined as the time occurring after the last time a close file change reason was entered in a change record, up to the time of the next succeeding close file change reason. For each file, an entry is made in the storage volume's master file table referencing the change record that currently includes the most up-to-date information on the status of the change reasons for the file. Deallocated change records are marked until a contiguous block is assembled, and then truncated from the change journal, returning the storage space to the file system for re-use.

Prior to truncation or deallocation, these change records are preferably archived by simply appending them to an archive change journal (step 110b) which, preferably, holds all previously deallocated change records. Thus, according to the present example embodiment, the archived change journal and the active change journal together contain all of the change records for the associate storage volume, providing a change history for each file created or modified since the change journal was created. This archival service would have the flexibility to store these records in any volume of its choice and to store any amount of them. (See, e.g., Abstract and column 11, lines 30-43).

However, unlike the present invention, Cabrera does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Cabrera at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach

or suggest these features in combination with the other limitations recited in each of the independent claims.

Wahl (U.S. Patent No. 6,324,654) discloses a computer network remote data mirroring system that writes update data both to a local data device and to a local, chronologically sequenced journal storage area, or writelog device. If the local computer system crashes, upon recovery or re-boot of the local computer system, the two most current updates in the writelog device are written to the local data device to assure that the data stored on the local data device is current. Additional memory or disk space is dynamically assigned to the writelog device to prevent a memory overflow condition. The computer network remote data mirroring system can be structured to provide logical groups of local data device/writelog device pairs. A primary mirror daemon on a local computer system monitors the writelog device for data updates and feeds the data over a network in the same order in which it is stored to a receiving remote mirror daemon on a remote computer system, which in turn commits the data updates to a mirror device. Failure recovery with the primary and remote mirror daemons is initiated automatically for certain failures which do not affect the basic operability of the overall computer system. A graphical user interface enables a user to configure the logical groups and create throttles, as well as to monitor performance of the remote data mirroring system. Network bandwidth throttling enables a predetermined portion of the network bandwidth to be assigned to remote data mirroring based on user-selected criteria. The writelog device is configured so that more memory space is dynamically assigned to the writelog

device to prevent a memory overflow condition which might otherwise corrupt stored data. In the embodiment in which the writelog device comprises a disk drive device, additional disk storage is dynamically assigned or another disk storage device is chained into the local, or primary, computer system. In the embodiment in which the writelog device comprises cache memory and a dirty bit map disk drive device, additional disk storage is dynamically assigned or another disk drive is chained into the local, or primary, computer system to prevent memory overflow. (See, e.g., Abstract and column 3, lines 9-20).

However, unlike the present invention, Wahl '654 does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Wahl '654 at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Wahl (U.S. Patent No. 6,442,706) discloses a computer network remote data mirroring system writes update data both to a local data device and to a local, chronologically sequenced journal storage area, or writelog device. A primary mirror daemon on a local computer system monitors the writelog device for data updates and feeds the data over a network in the same order in which it is stored to a receiving remote mirror daemon on a remote computer system, which in turn commits the data updates to a mirror device. A graphical user interface enables a user to create and configure throttles. Throttles are user-defined tests and actions evaluated by the primary mirror daemon to regulate network bandwidth, CPU, and writelog device utilization during data update mirroring. Network bandwidth throttling enables a predetermined portion of the network bandwidth to be assigned to remote data mirroring based on user-selected criteria. CPU throttling enables a user to control the amount of time the local data storage unit will wait prior to returning control to applications after an update. Writelog device throttling prevents a memory overflow condition by dynamically assigning memory to the writelog device by chaining writelog device extensions to the writelog device. Specifically, the writelog device is configured so that more memory or disk space is dynamically assigned to the writelog device to prevent a memory overflow condition which might otherwise corrupt stored data. In the embodiment in which the writelog device comprises a disk drive device, additional disk storage is dynamically assigned or another disk storage device is chained into the primary computer system. In the embodiment in which the writelog device comprises cache memory and a dirty bit map disk

drive device, additional disk storage is dynamically assigned or another disk storage device can be chained into the primary computer system to prevent memory overflow. (See, e.g., Abstract and column 7, line 65 to column 8, line 10).

However, unlike the present invention, Wahl '706 does not teach or suggest a backup and recovery method for a storage system that avoids journal overflow by, monitoring an amount of free space on the journal volume, and when the free space falls below a threshold value, taking a new snapshot (logical snapshot, stopping the storing of journal entries and switching to bitmap management, or overwriting the oldest journal entry) of the data volume and deleting the oldest journal entry, thereby avoiding journal overflow.

More particularly, Wahl '706 at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, the above described third feature of the present invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Therefore, since the cited references at a minimum fail to teach or the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 6, and the above described third feature of the present

invention as recited in independent claim 12, and the above described fourth feature of the present invention as recited in independent claim 19, and further fail to teach or suggest these features in combination with the other limitations recited in each of the independent claims, it is submitted that all of the claims are patentable over the cited references whether said references are taken individually or in combination with each other.

(F) Conclusion

Applicant has conducted what it believes to be a reasonable search, but makes no representation that "better" or more relevant prior art does not exist. The United States Patent and Trademark Office is urged to conduct its own complete search of the prior art, and to thoroughly examine this application in view of the prior art cited herein and any other prior art that the United States Patent and Trademark Office may locate in its own independent search. Further, while Applicant has identified in good faith certain portions of each of the references listed herein in order to provide the requisite detailed discussion of how the claimed subject matter is patentable over the references, the United States Patent and Trademark Office should not limit its review to the identified portions but rather, is urged to review and consider the entirety of each reference, and not to rely solely on the identified portions when examining this application.

In view of the foregoing, Applicant requests that this Petition to Make Special be granted and that the application undergo the accelerated examination procedure set forth in MPEP 708.02 VIII.

(G) Fee (37 C.F.R. 1.17(h))

The fee required by 37 C.F.R. § 1.17(h) is to be paid by:

☒ the Credit Card Payment Form (attached) for \$130.00.

☐ charging Account _____ the sum of \$130.00.

A duplicate of this petition is attached.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.,
Deposit Account No. 50-1417 (274.43199X00).

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.



Carl I. Brundidge
Reg. No. 29,621

CIB/jdc
(703) 684-1120